

What is claimed is:

1. A computer implemented method for group communication over a network of processors comprising:

determining an overlay spanning tree comprising an
5 origin node and at least one receiving node; and

controlling a source communication rate to be less than or equal to a bottleneck rate of the overlay spanning tree.

2. The computer implemented method of claim 1, further
10 comprising protecting data delivery by link error recovery.

3. The computer implemented method of claim 2, wherein the overlay spanning tree comprises a plurality of nodes, wherein the data delivery is reliable such that each node
15 receives the same data.

4. The computer implemented method of claim 1, further comprising scaling the overlay spanning tree to an arbitrary group size.

20

5. The computer implemented method of claim 1, further comprising determining a maximum throughput of the spanning tree among all possible configurations of the spanning tree given a reduced overlay distribution tree.

25

6. The computer implemented method of claim 5, wherein determining the overlay spanning tree comprises:

defining a target bandwidth for the overlay tree given a fully connected overlay distribution graph;

5 constructing a reduced overlay distribution graph by removing an edge from the fully connected overlay distribution graph having a bandwidth less than or equal to the target bandwidth;

constructing an arbitrary spanning tree comprising a
10 root, wherein the root is a source node of a plurality of links in the reduced overlay distribution graph;

performing a triangular improvement to remove a link violating a rate constraint;

increasing the target bandwidth upon determining that the
15 overlay spanning tree is constructible; and

decreasing the target bandwidth upon determining that the overlay spanning tree is not constructible.

7. The computer implemented method of claim 1, further
20 comprising joining a new node to the spanning tree.

8. The computer implemented method of claim 7, comprising joining the new node to an existing node of the spanning tree upon determining that the existing node has a bandwidth
25 of greater than or equal to an existing rate.

9. The computer implemented method of claim 8, further comprising:

determining a triangular improvement upon determining
5 that no existing node has a bandwidth greater than or equal to the existing rate;

joining the new node at an attachment point having a highest bandwidth among existing nodes of the spanning tree upon determining that the triangular improvement failed; and
10 redetermining the spanning tree upon determining bandwidth less than or equal to a minimum threshold.

10. The computer implemented method of claim 1, further comprising redetermining the spanning tree upon determining
15 that an existing node has left the spanning tree.

11. The computer implemented method of claim 10, further comprising:

determining orphaned child nodes of the existing node
20 that has left the spanning tree; and
performing a join for each orphaned child node.

12. A program storage device readable by machine, tangibly embodying a program of instructions executable by the

machine to perform method steps for group communication over a network of processors, the method steps comprising:

determining an overlay spanning tree comprising an origin node and at least one receiving node; and

5 controlling a source communication rate to be less than or equal to a bottleneck rate of the overlay spanning tree.

13. The method of claim 12, further comprising protecting data delivery by link error recovery.

10

14. The method of claim 13, wherein the overlay spanning tree comprises a plurality of nodes, wherein the data delivery is reliable such that each node receives the same data.

15

15. The method of claim 12, further comprising scaling the overlay spanning tree to an arbitrary group size.

16. The method of claim 12, further comprising determining
20 a maximum throughput of the spanning tree among all possible configurations of the spanning tree given a reduced overlay distribution tree.

17. The method of claim 16, wherein determining the overlay
25 spanning tree comprises:

defining a target bandwidth for the overlay tree given a fully connected overlay distribution graph;

constructing a reduced overlay distribution graph by removing an edge from the fully connected overlay

5 distribution graph having a bandwidth less than or equal to the target bandwidth;

constructing an arbitrary spanning tree comprising a root, wherein the root is a source node of a plurality of links in the reduced overlay distribution graph;

10 performing a triangular improvement to remove a link violating a rate constraint;

increasing the target bandwidth upon determining that the overlay spanning tree is constructible; and

15 decreasing the target bandwidth upon determining that the overlay spanning tree is not constructible.

18. The method of claim 12, further comprising joining a new node to the spanning tree.

20 19. The method of claim 18, comprising joining the new node to an existing node of the spanning tree upon determining that the existing node has a bandwidth of greater than or equal to an existing rate.

25 20. The method of claim 19, further comprising:

determining a triangular improvement upon determining that no existing node has a bandwidth greater than or equal to the existing rate;

5 joining the new node at an attachment point having a highest bandwidth among existing nodes of the spanning tree upon determining that the triangular improvement failed; and
redetermining the spanning tree upon determining bandwidth less than or equal to a minimum threshold.

10 21. The method of claim 12, further comprising redetermining the spanning tree upon determining that an existing node has left the spanning tree.

22. The method of claim 21, further comprising:
15 determining orphaned child nodes of the existing node that has left the spanning tree; and
performing a join for each orphaned child node.